From National Information Infrastructure
To Global Collaboration Infrastructure

Dr. Jean-Philippe Favreau and Dr. Kevin L. Mills
National Institute for Standards and Technology
United States of America

Abstract

Today, the information infrastructure within the United States provides individual users with wide access to a growing, diverse population of information sources. The emergence of Web technology attracts additional users onto the information infrastructure every day. As the population of cyberspace users increases, traditional forms of computer-based communication, such as electronic mail, become more widespread and users begin to experiment with newer forms, such as digital phones, chat groups, and video conferences. One day, soon, the national information infrastructure might be transformed into a global collaboration infrastructure. The essential underpinnings of such a transformation exist in the form of advanced research in networking, human-computer interaction, artificial intelligence, and digital libraries. To these underpinnings, we must add research in distributed collaboration technology, research that addresses the concept of groups and teams pursuing tasks and sharing information within virtual spaces while separated in time and in physical space. The successful outcome from such research could provide tomorrow's citizens with an electronic community through which computer-mediated collaboration can occur on a global scale. This paper describes the current state of the national information infrastructure, reviews some of the key technologies, both those subject to current research and those requiring new research, that could transform the national information infrastructure to a global collaboration infrastructure, and closes with a vision of tomorrow's electronic community.

1. Introduction

As we are moving toward the establishment of a new society -- the global information society -- it is critical to assess where does the future of information technology lie? In his book entitled _Future Telecommunications Information Applications, Services & Infrastructure_ [1], Robert K. Heldman contrasts his visions of two cities: Megatropolis and Infolopolis. Megatropolis bases its information infrastructure on islands of self-contained, private, and independent networks serving corporations or branches of the government. Public networks become de facto safety valves to handle overloads via specialized switches and gateways. In Megatropolis, face-to-face meetings continue as the best form of communication. Soon, Megatropolis experiences problems in exchanging information across communities of interests (e.g., "Doctors in one Health Management Organization (HMO) had difficulty sharing information with specialists or doctors in another HMO."); network addressing problems hamper cross-network accessibility, the complexity of the layers upon layers of protocols required to achieve interconnection creates bottlenecks and increases network congestion. Megatropolis's information infrastructure fails to satisfy the citizens increasing communication needs. In contrast, Infolopolis (distributed INFOrmation-based, Local megatroPOLIS) designs its communications infrastructure "to support new businesses that use information as one of their tools and assets."

Megatropolis and Infolopolis represent two contrasting visions from one man for the information infrastructure of our future. There are, of course, many different visions [2,3,6,7,8] of our future information infrastructure:

- a 500-channel interactive multimedia video/cable network; numerous "edutainment" multimedia products and services; telephone systems supporting voice, data, image, and video;
- an electronic marketplace for commercial and/or consumer products and services;
- a public network for government information and services, medical information, and education; or
- a source for discovering innovative applications for information technology in research, education, and commerce.

Common to all these visions, tomorrow's Global Information Infrastructure (GII) will offer a myriad of interconnected, interoperable information networks, computers, databases, and consumer appliances that will link homes, businesses, and branches of government together. These resources will provide citizens with universal network services, integration and translation services, data and knowledge management services, system software services, and will bring reliable computing and communications services to a growing, diverse population. Beyond support for the information needs of individual citizens, we might also envision enhanced capabilities for collaboration among groups and teams across the globe.

This paper presents the current state of the national information infrastructure, reviews some of the key technologies, both those subject to current research and those requiring new research, that could transform today's national information infrastructure to a global collaboration infrastructure, and closes with a vision of tomorrow's electronic community.

2. Current State of the NII

Used today by many different communities in support of collaboration, cooperation, and information dissemination, the Internet provides one key foundation of the NII. The Internet began in the late 1960's with the development of ARPANET, a 56-kbps backbone network. The University of California at Los Angeles installed the first ARPANET node in September 1969; by 1971 about 20 nodes were deployed. Developed from the mid-1970's to early 80's, the Transmission Control Protocol (TCP) and the Internet Protocol (IP) enabled networks of networks (so-called internetworks) to interoperate. From this beginning, the experimental ARPANET, initially limited to selected researchers who shared interactive communications between their computing systems at different locations, grew into today's world-wide Internet, a "network of networks" serving millions of users on every continent.

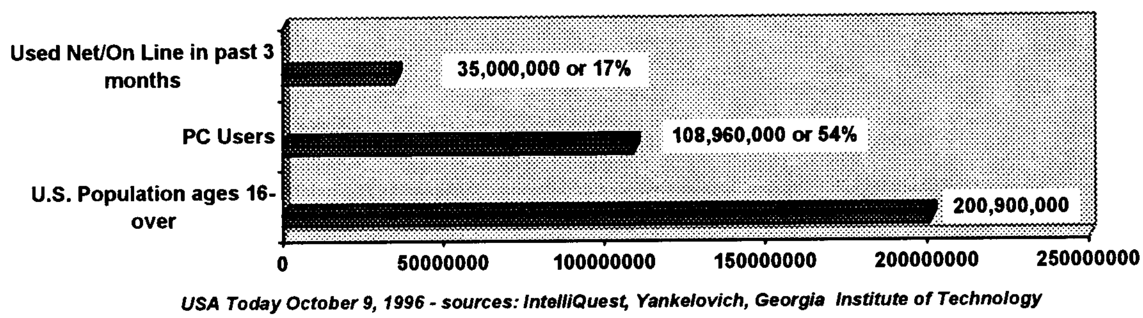

Figure 1: How many people are on-line?

The emergence of the World-Wide Web in the early 1990's and the development of Web browsers, such as Mosaic, transformed the way people view, exchange, and create information, and spawned a phenomenal growth of the Internet. The data and figures below illustrate the current state of the NII and GII, as viewed through the growth of the Internet in both the U.S. and world-wide:

- More than 35 million Americans now use the Internet -- 9 million of whom joined just this year (Cf. Figures 1, 2 and 3 [5]).
- Adults collectively spend as much time browsing the Internet each week as they do watching videocassettes. About 60 percent of Internet users are online for two hours or more per week. People

who use the Internet from work spend an average of eight hours per week online (Sources: FIND/SVP and Intelliquest, July 1996).

- More than 11.5 million adults use the World Wide Web [4].
- More than 1.5 million people have bought goods or services on the Web [4].

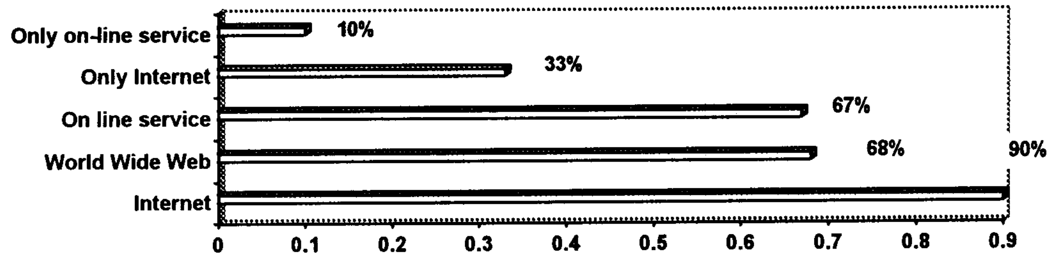

USA Today October 9, 1996 - sources: IntelliQuest, Yankelovich, Georgia Institute of Technology

Figure 2: Of the 35 millions adults, percentage who use:

- The InterNIC has registered nearly ½ million domain names. (source Internet Info, August 2, 1996)
- The July 1996 Domain Survey, which attempts to discover every host (uniquely reachable connected computer) on the Internet with a complete search of the Domain Name Server (DNS), uncovered 12.9 million hosts -- up from 9.5 million in January 1996. (Source: Network Wizards -- http://www.nw.com/, July 1996).

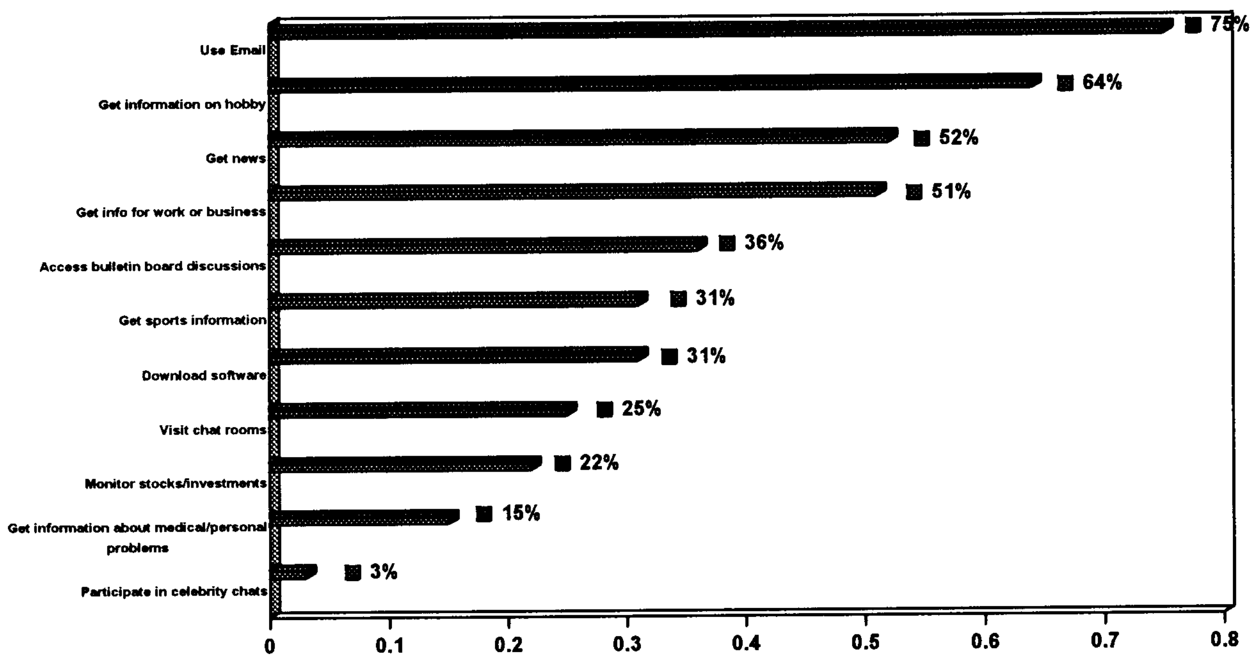

USA Today October 9, 1996 - sources: IntelliQuest, Yankelovich, Georgia Institute of Technology

Figure 3: On line activities

- The number of Internet hosts grew at an annual rate of 94 percent (Source: Network Wizards, July 1996).

- The global Internet market is expected to soar to $23 billion per year by 2000 (December 95 Source: Hambrecht & Quist).
- International Data Corp. predicts that by 1999 there will be 199 million Internet users and 38 million Web users.
- The Internet Society predicts that 120 million hosts will serve the Internet by end of this decade (January 1996 - The Internet Society).

3. Key Technologies

As the population of cyberspace users increases, traditional forms of computer-based communication, such as electronic mail, become more widespread and users begin to experiment with newer forms, such as digital phones, chat groups, and video conferences. One day, soon, the national information infrastructure might be transformed into a global collaboration infrastructure. The essential underpinnings of such a transformation exist in the form of advanced research in networking, human-computer interaction, artificial intelligence, and digital libraries.

3.1 Networking

Networking provides the underlying infrastructure that brings users together and enables information to move quickly, cheaply, and transparently across the globe. Today, electronic mail provides the fundamental network pipe; soon, multicasting services will enable streams of data to flow between groups of users; ultimately, high-speed network technology promises unlimited bandwidth available globally.

3.1.1 Electronic Mail - The Basic Communication Service For Today's GII

E-Mail technology started in the late 1960's with the ARPANET; use has now exploded with the rapid growth of the Internet (Cf. Figure 3). Starting as a means for small research groups to exchange information, E-Mail now provides the fundamental communication mechanism for the GII, [12] available to millions of people across the world and providing a powerful means to exchange information easily, quickly, and inexpensively in a variety of formats. E-Mail may be described as an asynchronous, electronic, interchange of information, supported with mechanisms allowing the creation, distribution, consumption, processing, and storage of messages [12]. Any individual can send a message to one or a group of individuals when the recipients are not present at their computer/appliance. The message, traveling at the speed of electricity or light through the networks, becomes available for viewing and reviewing at a time convenient for the recipient(s). While the technology started with the exchange of simple text messages, it is evolving into a powerful means of transmitting and viewing multimedia messages containing high-resolution color pictures, movie clips, and sound.

3.1.2 Multicasting and MBONE

The Internet's primary means of transmitting information moves individual packets that are each addressed to one site at a time (i.e., unicast). The next generation Internet will enhance the ability to send each packet of data to multiple sites at the same time (i.e., multicast) [10]. Multicast may be viewed as the Internet version of selective broadcasting in a very similar way to a television viewer selecting what is viewed by millions of other people. Multicast enables users to broadcast packets of information to anyone who is listening and has the right type of equipment and software.

Today, a small percentage of Internet routers are equipped to deal with multicast data streams. For this reason, an interim approach, called "multicast tunneling", has been developed to move multicast packets by putting them in regular unicast packets. The Internet is composed of islands, such as Local Area Networks (LANs), that can directly support IP multicast; these islands are connected by virtual point-to-point links, called "tunnels." Unicast packets flow through these tunnels only to be multicast on destination LANs. The first multicast tunnel was established between BBN and Stanford University in the summer of 1988. The IP Multicast Protocol (RFC 1112) was adopted by the Internet Engineering Task Force (IETF) in March of 1992 as the current, standard protocol for building multicast applications on the Internet.

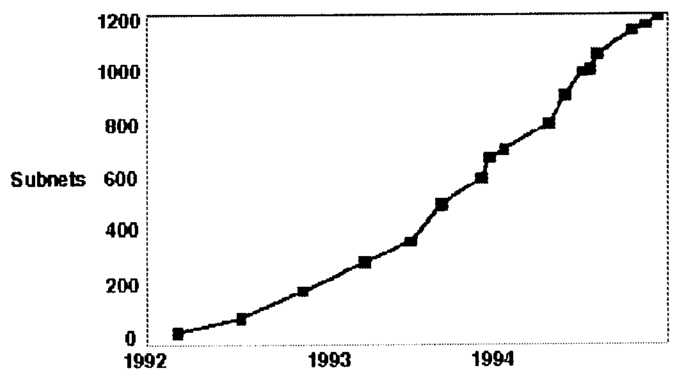

Figure 4: Growth in the MBone by year

The Virtual Internet Backbone for Multicast IP, or Mbone, is an experimental framework for developing and refining multicast protocols and applications on the Internet. The MBone is currently a co-operative voluntary effort, consisting of Internet service providers, who route multicast traffic over their networks, and end users who install multicast routers at their sites. While the MBone represents a tiny disjoint fragment of the entire Internet, it has experienced exponential growth in the number of participating sites since its inception in 1992 (Cf. Figure 4) and now spans over several continents. The current growth in popularity of multimedia collaborative applications and the support for native multicasting expected in the next generation Internet guarantees an important role for IP Multicast in the future.

Table 1: ATDnet

The Advanced Technology Demonstration Network (ATDnet) is a recently inaugurated high performance networking testbed in the Washington D.C. area. It is intended to be representative of possible future Metropolitan Area Networks. Established by the Advanced Research Projects Agency (ARPA) to enable collaboration among Defense and other Federal agencies, ATDnet has a primary goal to serve as an experimental platform for diverse network research and demonstration initiatives. Emphasis is on early deployment of emerging Asynchronous Transfer Mode (ATM) and Synchronous Optical Network (SONET) technologies. "ATM/SONET" refers to transmitting ATM packets over a Synchronous Optical Network (SONET) link, a standard for transmitting data across fiber optic cable.

The ATDnet concept is to interconnect several agency sites with high speed fiber optic transmission media and to overlay this media with SONET and ATM protocols. The initial deployment will be at OC-48 data rates (approximately 2.5 billion bits/second, also referred to as 2.5 gigabits/second) but is designed to scale upwards to technology-limited data rates. Pair-wise and multiple party research initiatives and experiments are planned over the lifetime of the testbed. Experimental testing of the bitways and a diversity of service and applications experiments are intended to gain insight into the potential of this new performance level.

In FY 1995 ATDnet was used for experiments in telemedicine, distributed simulation, security and encryption techniques, ubiquitous video teleconferencing, large ATM network signaling research, ship design and analysis visualization, and network experiments. New network management techniques for SONET and ATM systems are being developed that will give DOD fast-start access to equivalent commercial services once they are tariffed.

Source ATDNet (http://www.atd.net/atdnet.html)

3.1.3 High Performance Networks - Asynchronous Transfer Mode

Emerging as the primary network technology for next-generation, multimedia communication, the Asynchronous Transfer Mode (ATM) was born from standardization efforts for Broadband Integrated Services Digital Network, (B-ISDN) which began in the CCITT in the mid 1980s. ATM provides a specific connection-oriented packet switching transfer mode based on fixed length cells (i.e., 53 bytes - 5 bytes for header and 48 bytes for information) [15]. ATM cells flow transparently through the network, such that no error control processing protects the information field of ATM cells inside the network. ATM protocols are designed to handle isochronous (time critical) data such as video and telephony (audio), in addition to more conventional data communications between computers.

ATM was one of the network technologies deployed in the gigabit testbeds (see example in Table 1) funded jointly by the National Science Foundation (NSF) and the Defense Advanced Research Projects Agency (DARPA) as part of the U.S. Government's High Performance Computing and Communications (HPCC) program. These gigabit testbeds provided the means to test applications and extended the state of the art in high-speed networking. The accomplishments of the gigabit testbed projects include:

- Operation of high speed data transmission systems not previously achieved in long distance networks (2.4 billion bits per second across thousands of kilometers).
- New very high-speed transport protocols for wide-area networks (new transmission records of 500-800 Mb/s between hosts over long distances)
- New ATM switches to route data at higher speed across the testbeds (pushed switching to 155 Mb/s, 622 Mb/s, and faster).

Table 2: Major elements of HCI research

- New interaction paradigms that enable users to communicate with computers and with others using computers in more effective, human ways;
- Fundamental understanding of the nature of interaction that will enable users to achieve their goals more directly, including a better understanding of the limitations of human information processing; and
- Engineering methods for interactive system design and development that will improve the quality and reduce the costs of user interfaces for new systems.

Source: New Directions in Human-Computer Interaction Education, Research and Practice [16]

3.2. Human-Computer Interaction

Fast changing network and computing technologies are affecting the way we conduct business. Key to our ability to absorb new information and to apply new technologies is the means by which we interact with computer systems and networks. Effective human-computer interaction can improve the ability of the ordinary citizen to use information technology and to access information; alternatively, ineffective human-computer interaction can create fundamental barriers to reaching the promises of the NII. [13, 16].

> *"If the interface is ineffective, the system's functionality and usefulness are limited; users become confused, frustrated, and annoyed; developers lose credibility; and the organization is saddled with high support costs and low productivity." [17]*

Human-Computer Interaction (HCI) is the study of how people design, implement, and use interactive computer systems, and how computers affect individuals, organizations, and society. The reasons for our difficulties in interacting with computer interfaces are multiple, including a misunderstanding of the user needs and of the tasks to be performed, a misunderstanding of human cognitive abilities, and a lack of guidelines and metrics to guide the development of advanced user interfaces. Improvements in system usability will require the identification and promotion of technologies addressing these problems (Cf. Table 2).

The development of technologies for HCI can be seen as a process of adapting computers and networks to humans. Key research bridging the gap between information systems and humans includes technologies that exploit modalities used naturally by humans. New technologies related to speech, handwriting, natural language, and other modes of interaction will have a lasting impact upon the work environment, and will open up new areas of application for information technology [13,20]. For example, the Virtual Reality Modeling Language (VRML) constitutes a clear evidence of the migration from two-dimensional presentation to three dimensional space. The notions of selecting and interacting with information will need to be revised, and techniques for navigation through information spaces will need to be radically altered from the present two-dimensional page-based models.

3.3 Artificial Intelligence

The exponential growth of the Internet has produced an overload of information including electronic documents, databases, and services. Bringing dispersed library resources and entertainment materials to the ordinary citizen leads to wasted time and effort searching for the appropriate sources of information. Furthermore, as the information is dispersed over several sources, the user must go through the process of accessing several information sources, learning their different user interfaces and query languages, and, later, undertaking the cumbersome task of information integration.

Table 3: Present-day division of AI Research

Knowledge representation and articulation, which seeks to discover expressive and efficient forms and methods for representing information about all aspects of the world and to use these methods to create and compile explicit, formal, multipurpose catalogs of substantive knowledge;

Learning and adaptation, which extends statistical, analytical, and scientific discovery techniques and hypothesized neurophysiological mechanisms to procedures that extract a wide range of general trends, facts, and techniques from instruction, experience, and collected data;

Deliberation, planning, and acting, which concerns methods for making decisions, constructing plans or designs to achieve specified goals, and monitoring, interpreting, diagnosing, and correcting the performance of the plans and implementations of the designs.

Speech and language processing, which seeks to create systems capable of communicating in and translating among natural written and spoken languages;

Image understanding and synthesis, which develops algorithms for analyzing photographs, diagrams, and real-time video image streams as well as techniques for the customized presentation of quantitative and structured information;

Manipulation and locomotion, which seeks to replicate and surpass the abilities of natural hands, arms, feet, and bodies;

Autonomous agents and robots, which integrates the other areas to create robust, active entities capable of independent, intelligent, real-time interactions with an environment over an extended period;

Multi-agent systems, which identify the knowledge, representations, and procedures needed by agents to work together or around each other;

Cognitive modeling, which focuses on contributing techniques and constructing integrated architectures that replicate structural or behavioral features of human cognition; and

Mathematical foundations, which takes the concepts and techniques of the other areas as subjects for formalization, distillation, analysis, and re-conceptualization.

Source: Strategic Directions in Artificial Intelligence [14]

For the information to be accessible and for the NII to be successful, the ordinary citizen will need the necessary tools and technologies to navigate a labyrinth of databases and services. The field of artificial intelligence (Cf. Table 3) can play a pivotal role in meeting these challenges [11,14] in technologies such as:

- goal-oriented, cooperative, and customizable human-computer interfaces that handle different modalities (e.g., natural language, gesture, graphics, or animation), operate as intelligent systems that interpret the requests from the users and perform the required tasks, and adjust themselves automatically to specific users;
- intelligent indexing, access, and retrieval of all forms of information; and
- development of software tools and environments to support all stages of a project's life cycle -- specification, design, adaptation, construction, evaluation, and maintenance.

3.4 Digital Libraries

The phenomenal growth of the Internet and the World-Wide Web is transforming the traditional role of libraries. Indeed, the carefully selected, acquired, and cataloged collections of items stored within the facilities of the "classical" library are only available to the individuals that physically enter the library building. The "Library of the NII" constitutes a source of inexpensive and extensive electronic collections of digital information -- including text, pictures, audio, and video -- accessible remotely and simultaneously to readers located at multiple sites [9]. These Digital Libraries store materials in electronic format and offer effective manipulation (search and access) of those materials.

While research into digital libraries is intimately related to research into distributed information systems (i.e., to effectively mass-manipulate the information on the Internet), the most challenging activities are targeted towards "enabling technology and infrastructure for wide-area information management, exchange, and collaboration" [18].

Table 4: NSF/DARPA/NASA Funded Digital Libraries Initiative

University of California, Berkeley: broad-range project including image analysis, natural language analysis, and computer vision analysis for effective information extraction; new user interface paradigms and authoring tools for better accessing of multimedia information; and improved protocols for client program interaction with repositories.

University of California, Santa Barbara: to explore a variety of research and development problems related to a distributed digital library for geographically-referenced information. The initial objective is to turn a large cartographic collection -- now totaling more than 4 million items -- into electronically accessible forms, and to build a system that allows users to locate maps and other spatially-indexed materials in geographically dispersed digital libraries and databases.

Carnegie Mellon University: to establish a large, on-line digital video library by developing intelligent, automatic mechanisms to populate the library and allow for full-content and knowledge-based search and retrieval via desktop computer and metropolitan area networks.

University of Michigan: to create an infrastructure for rendering library services over a digital network. This project will conduct coordinated research and development to create, operate, use and evaluate a testbed of a large-scale, continually evolving multimedia digital library. The content focus of the library will be earth and space sciences.

University of Illinois at Urbana-Champaign: to build repositories of indexed multiple-source collections and federate them by searching the material via multiple views of a single virtual collection. Prototypes will be built to make library information readily accessible to Internet browsers by forming the "Interspace" in which information is linked coherently and displayed in appropriate formats.

Stanford University: To address the problem of interoperability and to develop the enabling technologies for a single integrated "virtual" library capable of providing uniform access to a large number of emerging networked information sources and collections. Digital Library testbed focuses on interoperability among existing collections and publication-related services.

Source: Computer, theme issue on the US Digital Library Initiative, May 1996 [18]

The development of digital collections in libraries will depend on several components/technologies including interconnected and interoperable networks, decentralized data and processing, databases, navigation and retrieval tools, document delivery, presentation standards and techniques, mass storage, and human resources. Research in areas related to storing, finding, transmitting, viewing, manipulating, and controlling access to complex information are keys to this core element of the NII. In 1994 the NSF, DARPA, and the National Aeronautics and Space Administration (NASA) jointly funded a four-year digital libraries research and technology development activity. These projects, briefly introduced in Table 4, are in the process of developing the next generation of tools for information discovery, management, retrieval and analysis [18,19].

4. Challenges For Collaboration Researchers

While current research in networks, human-computer interaction, artificial intelligence, and digital libraries promises to improve the ability of individual users to exploit the evolving GII, to realize the potential of a global collaboration infrastructure, new research challenges must be met in at least four areas: 1) collaboration middleware, 2) tools for sharing meaning, 3) tools for sharing views, and 4) evaluation metrics and methodologies. Collaboration middleware refers to the software that binds collaborators and information together, either synchronously or asynchronously, through multiple multi-modal channels across time and space. Tools for sharing meaning encompass software components that place collaborators and information within the context of a shared task and that enable collaborators to overcome differences in terminology and language. Tools for sharing views comprise software that enables collaborators to interact through correlated visualizations that represent the information space pertaining to a collaborative task. Evaluation metrics and methodologies provide interesting challenges because users need a means to distinguish technologies that improve their abilities to collaborate from those that do not and because researchers need a means to assess their technical approaches against significant, measurable metrics. The following sections address some specific research challenges in each of these four areas.

4.1 Collaboration Middleware

In its current state, collaboration software comes in diverse forms, including audio and video conferencing tools, client-server systems for voting and for discourse management, multi-user domains (MUDs), shared whiteboards, work flow management systems, and shared, special purpose applications. Each instance of these forms provides software to perform specific functions and enables operation on a limited set of hardware and software platforms. Integrating various collaboration services together presents developers with a difficult challenge, and achieving interoperation between different implementations of similar services often proves even more difficult. Achieving such integration and interoperability within a software framework that enables local and global resources to be managed effectively remains beyond the state of the art. A need exists for middleware that can exploit the underlying services of networks and distributed operating systems to provide collaboration services to users and to application software. Providing such middleware requires solutions to three significant research problems.

Challenge #1 Scalable, Reliable Multicast Services

Synchronous collaborative sessions generally involve more than two individuals; thus, session data must be routed to multiple, physically separated locations. For loss-tolerant data streams, such as digital video and audio, several protocols exist that allow data to be multicast efficiently to all end points within a session. In general, these protocols rely on receiving nodes to detect and recover from packet losses without additional interactions on the multi-party channel. For loss-intolerant data streams, such as whiteboard transactions and distributed animation events, no effective and efficient multicast protocols exist. Research is needed to develop, test, and evaluate scalable protocols that enable data to be multicast reliably to multiple parties. Here, scalable means that error recovery can be effected with a minimum of overhead transmissions on a multi-party channel.

Challenge #2 Resource Management

Collaborative applications present complex resource management requirements well beyond those faced by less demanding distributed systems. At each node in a collaboration session, local session management algorithms are required to: 1) arbitrate access to non-shareable devices, 2) allocate local resources among modalities (such as voice, text, and video) based on a user's varying task focus, and 3) synchronize timing among related but separate data streams. Across all nodes and links comprising a session, global session management algorithms are required to: 1) adaptively allocate distributed bandwidth and computation among competing demands within a session, 2) manage distributed control of the floor, and 3) regulate distributed access to shared conference resources (such as video cameras and microphones within a conference room). In heterogeneous environments, resource management techniques must identify the need for translating encodings (i.e., transcoding) among various video and audio formats, control protocols, and display formats. As an ideal, resource managers might also detect the need for various forms of transcoding within a session, arrange for appropriate transcoders to be generated, and then select an appropriate mapping of transcoders to system resources based upon combinations of resource availability and requirements for performance, function, and security. Given the dynamic nature of most collaborative sessions, transcoders should also be considered within the adaptive resource allocation framework for global session management.

Challenge #3 Composable and Interoperable Architectures

Generally, little understanding exists regarding the components and relationships required to realize a comprehensive and flexible architecture for collaboration systems. Development of such architectures could lead to numerous benefits. Existing, stand-alone applications could be made collaborative with little effort. For example, given a sophisticated, single-user virtual reality walk-through model of a building design, a team of people could collaboratively walk-through a virtual building. Users could interact through a collaborative space using familiar tools yet still achieve interoperable semantic operations. For example, given two groups of users collaborating to write a document, each user could edit the document using the text editor most comfortable to him, rather than having some set of users (or worse all users) switch to an unfamiliar, group-enabled, text editor. Improved collaboration algorithms and components could be placed into existing software systems with minimal adverse effects. For example, as improved protocols are developed for synchronizing audio, video, and timed event streams, implementations of the new algorithms could be substituted into existing collaboration software with minimal changes to other system components or to collaboration client software.

4.2 Tools for Sharing Meaning

In their current state, most successful collaboration products enable collaborators to exchange representations of ideas encoded for computer transmission and rendering. For example, faxes, video and audio streams, and text files can all be used to encode and then convey representations of concepts. The context required to interpret the encoded concepts and to relate them to other, similarly encoded, concepts remains outside the state-of-the-art for most collaboration products. Such products lack suitable representations for semantic constructs that are: 1) meaningful to human users and that can also be 2) processed by computer software. This lack of support for semantics plagues both the application information exchanged as well as the collaboration process itself. Even the simplest concepts such as the identities and roles of collaborators and the specific agenda or process being followed are not typically represented, nor are any relationships among these concepts or between these concepts and relevant information artifacts pertaining to a collaboration. Nowhere are these deficiencies more apparent than in the simple lack of support in most collaboration software for effective archiving. Without support for semantics surrounding the goals, objectives, and activities of collaborators, current collaboration systems also require that a user know the identity of fellow collaborators and that a user identify the information to be used during a collaboration -- without semantics, the system is in no position to assist in the discovery of potentially fruitful collaborators or in the discovery of information relevant to the collaboration. While most commercially successful collaboration products provide no support for the collaboration process, instead leaving collaborators to devise a process completely outside the ken of the collaboration software, some less successful collaboration products provide process support in the form of programmer-encoded workflows that often prove too restrictive, yet also too difficult for the user to adjust. These serious semantic deficiencies suggest some potentially profitable research directions.

Challenge #4 Incremental Semantic Models

Groups and teams collaborate to solve a problem or to achieve a common objective. In most cases, the collaboration process requires that a group develop, over the course of time, a shared understanding of the semantics of the information related to their problem or objective. Most existing modeling tools provide either no support for semantic description (relying on human interpretation to impart semantics) or rigorous support for semantic description (requiring full specification of a knowledge base). Unfortunately, most collaborations operate somewhere between these two extremes. For this reason, effective collaboration systems must support the incremental, evolutionary development of shared semantic models. Such semantic models must permit incompleteness, inconsistency, and ambiguity. In addition, facilities must be provided to enable collaborators to merge separately developed semantic models into shared models. Semantic modeling facilities must also enable collaborators to increase the completeness, consistency, and precision of a shared model as collaborators discover new information and achieve consensus.

Challenge #5 Semantic Indexing Of Timed Events and Multimedia Streams

Today, most of the deliberations and interactions within collaborative systems are simply lost to those not actually present. This fact is unfortunate for those who enter a collaboration late and wish to catch-up, and for those who wish to conduct a retrospective review of what happened and why. Assuming that disk and data capture technology advance to the point where it becomes feasible to record the entire history of a collaboration, how can users access the resulting immense archive of material? All event and multimedia streams within the archive must, at minimum, be time tagged. Such an archive would enable a user to browse the archive based on only one question: What was happening at time T? More interesting questions (such as: What did my competitors say? and When were new collaboration products discussed?) cannot be answered. Methods must be found to semantically index collaboration archives so that users can retrieve relevant aspects of the collaboration based on specific interests. Such methods might involve recording the time tags of semantically significant events in the archive, either in real-time during a collaborative session or off-line through post-processing. An official archive index might reflect someone's semantic view of what events would interest people later. Value-added indexes might be created, using post-processing, to reflect the time tags of significant events related to questions that specific users want to answer. Multiple informal archive indexes might be generated by individual collaborators who register specific time tags within their personal notes. Such personal notes would than contain a particular semantic view over the collaboration archive. Using such an approach, multiple collaborators might share their personal views with each other, much as students exchange class notes.

Challenge #6 Semantic Models For Representing Goals, Objectives, and Tasks

Today's collaboration systems offer users little significant help to find potential collaborators or relevant information, nor to relate the actions and artifacts of others to a user's specific goals, objectives, and tasks. No progress can be expected in this area until computer systems are able to represent the semantic context of a user's activities (e.g., goals, objectives, tasks and the relationships between them) and to relate a user's actions and information to that semantic context. Research is required regarding methods for representing reusable and composable semantic components, and semantic frameworks in which such components can be employed. Given the availability of such semantic frameworks, research is needed to relate a user's actions within a semantic context and to provide methods for comparing and contrasting a user's semantic context with the semantic contexts of a large population of potential collaborators and information sources. Once collaboration software has access to descriptions of what users are trying to accomplish and why, intelligent software agents can begin to provide value-added assistance to collaborators.

Challenge #7 Group-Evolvable Processes

Rarely does one process suit all collaborating groups. For example, the brainstorming process suitable for a marketing group would not prove suitable for a command and control team monitoring and adjusting the execution of a military mission. Further, rarely does one process suit a given collaborating group for all time. For example, a collaboration that started out as an informal gathering of four people with similar interests might evolve to a professional society with thousands of members. Effective processes for four like-minded individuals are unlikely to prove productive in a large organization. From these examples, one can see that computer support for collaborative processes must be flexible and evolvable. The ultimate in flexibility occurs when no process constraints are enforced by a collaboration system. This approach works well until the first, second, or third situation where the collaborators would benefit from support for process. Once collaborators need support for process most collaboration systems require someone to work out the required constraints and then to encode them in a computer language. Once encoded, the process constraints can quickly become too constraining. Research is required to explore methods for groups to gracefully define and evolve process constraints in response to the realistic needs of the collaborators. Once a group realizes some constraints are necessary, the group should be able to propose, examine, test, and then jointly decide to impose the constraints. A truly sage collaboration system, containing knowledge about effective processes for particular tasks, team types, and group sizes, might suggest process constraints based upon monitoring of a team's interactions and progress.

Challenge #8 Semantic Interoperability

Potentially productive collaborations, among individuals from different backgrounds, often founder on the inability of collaborators to understand one another. Two particular problems arise. In the first case, individuals might possess different semantic models described using overlapping symbols. For example, the networking expert's IP (Internet Protocol) maps syntactically but not semantically to the VLSI designer's IP (Intellectual Property). In the second case, individuals speaking only French will encounter difficulty collaborating with individuals speaking only Greek. From these examples, one can see that research must address the problems of semantic mapping among concept spaces in multiple domains and among natural languages. In synchronous collaboration settings, the semantic mapping problem, difficult in any case, increases due to requirements to process spoken natural language and to provide real-time performance.

4.3 Tools for Sharing Views

For most commercial collaboration products sharing views means sharing a whiteboard -- a basic bit map onto which users can draw other bit maps. More usefully, most products also support a shared audio channel that enables a group of users to explain themselves. This level of support for shared views, while universal, disappoints. What support for shared views might prove more inspiring? Suppose collaboration systems could provide interconnected, role-specific and task-specific views driven from a shared underlying domain model. Within such a system, a military commander might share a basic geographic view with his staff, while, bounded by a geographic area designated by the commander, the logistics officer would see supply lines, depots, and inventories, the intelligence officer would see suspected enemy troop deployments, and the meteorology officer would see relevant weather information. From such a powerful starting point, where each officer views information relevant to his special expertise but also tied to a common view, the commander and staff could discuss upcoming battle options and each specialist could share information onto the common view as relevant within the context of the discussion. Suppose also that someone enters the discussion later in time and desires to understand the foregoing deliberations. How might this goal be met? Perhaps the shared views could be annotated, using multimedia techniques, to capture the rationale for decisions made or the concerns discussed at key points. Suppose that a commander's battle plan could be animated and that such an animation could be shared among a group, where each member can stop the animation, backup, move forward, and annotate particular events of interests during a collaborative discussion. Current products can provide little of the potential power that shared views can bring to collaborations. Visualization research exists today but some challenges derive directly from sharing visualizations among groups.

Challenge #9 Control Protocols For Distributed View Sharing

Visualization researchers today pursue sophisticated methods for adaptively representing information in forms that match a display context or even a user's goals. Such research needs to continue; however, additional wrinkles must be considered when teams can take multiple views into a shared information space. Assuming collaborating users can develop independent views, how can groups be made to see a visual effect from the viewpoint of a specific group member? When each user within a group must take a particular, role-defined, logical view of the information space, how are changes made in one member's view logically reflected in the views of other members? How can access to a shared viewing space be controlled?

Challenge #10 Visualizing Abstract Spaces

Many methods exist for visualizing data that maps nicely to 2-D, 3-D, and sometimes even 4-D spaces. Difficulties arise when more abstract information must be visualized. Some current research focuses on methods for visualizing abstract information within the context of 3-D and 4-D views. For example, message traffic statistics can be mapped to a geographic map, vital signs and other medical measurements can be mapped to appropriate images of the human body, and design specifications can be mapped to relevant components within a physical design rendering. Further research along such lines should lead to methods for augmenting real spaces with superimposed information that adds value. For example, imagine that a soldier looks at an approaching aircraft and can see the speed, altitude, heading, aircraft type, and projected flight path. More difficult problems face researchers seeking techniques to visualize abstract (N-dimensional) spaces. For example, how should an executing distributed software system appear? What does a half-executed military plan look like? In cases such as these, no intuitive visual mapping appears obvious. Further, navigating such abstract spaces can be difficult because users cannot find intuitive anchors to guide them. Research is needed to develop methods for representing abstract information spaces and for navigating such spaces. Perhaps multimedia annotations can be inserted within abstract representations in order to provide users with navigation tracks and to explain the significance of the visual representation at particular points in N-space and time.

4.4 Evaluation Metrics and Methodologies

For researchers developing advanced collaboration technologies, methods are needed to assess the effectiveness and efficiency of proposed solutions. To date, few widely agreed metrics or methodologies exist to help collaboration researchers evaluate their technologies against the needs of users. In general, user objectives are task-based; thus, even where user objectives are defined, those objectives must be decomposed into relevant metrics that can be measured.

Challenge #11 Measurement

Given the scope of information technology required to support collaborations, at least three classes of measures might prove useful: 1) measures of the efficiency and sufficiency of the underlying technology connecting users and information together, 2) measures of the effectiveness of various human-system interaction techniques, and 3) measures of the productivity effects realized from using a collaboration system to accomplish tasks. While some significant research experience exists with the first two of these classes, the third class remains less well understood. Perhaps researchers need to abstract meaningful task and team models for collaborative activities, then to devise salient measures for the selected models, and finally to define metrics that can be used to express the salient measures. Given a clear definition for metrics, what components of a system should be instrumented, and how can such instrumentation be effected with minimal distortion to the instrumented system? How can measurement experiments be made repeatable? How can experiments be automated to the maximum extent possible? How can measurements requiring significant system load be obtained in laboratory settings?

5. Virtual Electronic Communities

Once the current research in networking, human-computer interaction, artificial intelligence, and digital libraries matures and after researchers meet the dozen difficult challenges impeding the use of the GII as a medium for collaboration, the citizens of Infolopolis might expand their horizons to encompass participation in and contribution to a wide array of virtual electronic communities. Imagine Infolopolan Ingmar Prozeman flipping on her computer and instantly completing her morning commute to a virtual meeting room where she consults with fellow members of a team designing a new electronic shopping mall. At noon, Ingmar switches virtually to a conference with her daughter's teachers who themselves are actually spreadout across the globe. In the afternoon, Ingmar joins a virtual panel of fellow members of her where she discusses and evaluates candidates for this year's awards. Before joining her daughter to inspect progress on several school assignments arrayed across the computers comprising her daughter's virtual classroom, Ingmar "drops by" the design team meeting room to see if anyone is still working and to check on progress since her earlier visit that morning. Later in the afternoon, she "makes" a virtual visit to her mother on the West coast, chats a bit, "pulls" the family's last vacation video-clip from her video bank account and shares it with her. During dinner, the family discusses their plan to visit Ingmar's sister abroad. While her husband switches the television to the Travel Services Network, Ingmar establishes a virtual call to her sister, and starts discussing the most convenient dates while her husband orally requests information about the best itinerary, and reserves a hotel and transportation. After dinner, Ingmar visits her virtual bookclub where several members present some ideas for books to read over the next few months.

Bibliography

[1] Heldman Robert, K., *Future Telecommunications Information Applications, Services, and Infrastructure*, McGraw-Hill, 1993.

[2] *Revolution in the U.S. Information Structure*, National Research Council , National Academy Press, June 1994 http://www.nap.edu/readingroom/books/newpath/.

[3] *Realizing the Information Future: The Internet and Beyond*, National Research Council , National Academy Press, June 1994 - http://www.nap.edu/readingroom/books/rtif.

[4] D. Hoffman and T. Novak, *Internet and Web Use in the United States: Baselines for Commercial Development*, Owen Graduate School of Management, Vanderbilt University, Nashville, TN 37203, July 1996 - http://www2000.ogsm.vanderbilt.edu/.

[5] "On-Line Elusive Face," USA Today, October 9, 1996.

[6] *Evolving the High Performance Computing and Communications Initiative to Support the Nation's Information Infrastructure*, National Research Council, National Academy Press, June 1995 - http://www.nap.edu/readingroom/books/hpcc.

[7] *The Unpredictable Certainty Information Infrastructure Through 2000*, National Research Council , National Academy Press, Washington, D.C. 1996 http://www.nap.edu/readingroom/books/unpredictable/.

[8] *Framework for NII Services*, National Institute of Standards and Technology, NISTIR 5478, 1994 http://dsys.ncsl.nist.gov/pub/sawg/niiframe.html.

[9] *High Performance Computing and Communications: Foundation for America's Information Future*, The High Performance Computing and Communication (HPCC) Blue Book, 1996 - http://www.hpcc.gov/blue96/index.html.

[10] Kevin Savetz, Neil Randall, and Yves Lepage , *MBONE: Multicasting Tomorrow's Internet*, IDG Books Worldwide, Inc., April, 1996

[11] *The Role of Intelligent Systems in the National Information Infrastructure*, The American Association for Artificial Intelligence, Edited by Daniel S. Weld, University of Washington - http://www.aaai.org:80/Policy/Papers/nii.html.

[12] R.H. Anderson, T.K. Bikson, S.A Law, Bridger M. Mitchell, *Universal access to E-mail, Feasibility and Societal Implications*, MR-650-MF, Published in 1995 by RAND - http://www.rand.org/publications/MR/MR650/.

[13] "Strategic Directions in Human-Computer Interaction", (*This article is based on the results of the Human Computer Interaction Working Group of the ACM Workshop on Strategic Directions in Computing Research, September 1996 - http://www.cs.brown.edu/people/ifc/hci.html, Draft of an article to appear in ACM Computing Surveys 28(4), December 1996.

[14] "Strategic Directions in Artificial Intelligence", (*This article is based on the results of Artificial Intelligence Working Group of the ACM Workshop on Strategic Directions in Computing Research, September 1996 - http://www.medg.lcs.mit.edu/doyle/sdcr/ai/report/latest-draft.html. Draft of an article to appear in ACM Computing Surveys 28(4), December 1996.

[15] D. E. McDysan; D. L. Spohn, ATM Theory and Application, Mc-Graw-Hill, Series on Computer Communications, 1994.

[16] G.W. Strong, A Report: NSF Workshop on New Directions in Human-Computer Interaction Education, Research and Practice. September, 1994 - http://www.sei.cmu.edu/arpa/hci/directions/TitlePage.html.

[17] Nolan, Norton and Co. *Managing End-User Computing.* Boston: Nolan, Norton & Co., 1992.

[18] *Computer,* theme issue on the US Digital Library Initiative, May 1996 - http://www.computer.org/pubs/computer/dli.

[19] "The DLI Testbeds: Today and Tomorrow," *D-Lib Magazine,* July/August 1996. - http://www.dlib.org.

[20] A *Survey of the State of the Art in Human Language Technology*, Center for Spoken Language Understanding, Oregon Graduate Institute of Science and Technology, January 1996 - http://www.cse.ogi.edu/CSLU/HLTsurvey.